

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	70	(logical adj disk) with parity	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:09
L2	1897	711/114.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:10
L3	70	(logical adj disk) with parity	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:10
L4	1897	711/114.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:10
L5	44	L3 and L4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:10
L6	2188	logical adj entit\$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:11
L7	44	5 and parity	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:12
L8	16022	load adj balanc\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:13
L9	9487884	(load adj balanc\$4) or ((High or more) access)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:16

L10	44	5 and 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:14
L11	10092812	(load adj balanc\$4) or ((High\$3 or more) access) or (load adj leveling)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:16
L12	44	5 and 11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 12:16


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

raid and (logical disk) and (logical entities) and (load balancing)



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

raid and logical disk and logical entities and load balancing

Found 35,303 of 166,357

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Run-time adaptation in river](#)



Remzi H. Arpaci-Dusseau

February 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 1

Publisher: ACM Press

Full text available: [pdf\(849.04 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of run-time adaptation within the River dataflow programming environment. The goal of the River system is to provide adaptive mechanisms that allow database query-processing applications to cope with performance variations that are common in cluster platforms. We describe the system and its basic mechanisms, and carefully evaluate those mechanisms and their effectiveness. In our analysis, we answer four previously unanswered and important que ...

Keywords: Performance availability, clusters, parallel I/O, performance faults, robust performance, run-time adaptation

2 [I/O reference behavior of production database workloads and the TPC benchmarks— an analysis at the logical level](#)



Windsor W. Hsu, Alan Jay Smith, Honesty C. Young

March 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 1

Publisher: ACM Press

Full text available: [pdf\(5.42 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As improvements in processor performance continue to far outpace improvements in storage performance, I/O is increasingly the bottleneck in computer systems, especially in large database systems that manage huge amounts of data. The key to achieving good I/O performance is to thoroughly understand its characteristics. In this article we present a comprehensive analysis of the logical I/O reference behavior of the peak production database workloads from ten of the world's largest corporatio ...


Keywords: I/O, TPC benchmarks, caching, locality, prefetching, production database workloads, reference behavior, sequentiality, workload characterization

3 [Minerva: An automated resource provisioning tool for large-scale storage systems](#)



-  Guillermo A. Alvarez, Elizabeth Borowsky, Susie Go, Theodore H. Romer, Ralph Becker-Szendy, Richard Golding, Arif Merchant, Mirjana Spasojevic, Alistair Veitch, John Wilkes
November 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(701.98 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Enterprise-scale storage systems, which can contain hundreds of host computers and storage devices and up to tens of thousands of disks and logical volumes, are difficult to design. The volume of choices that need to be made is massive, and many choices have unforeseen interactions. Storage system design is tedious and complicated to do by hand, usually leading to solutions that are grossly over-provisioned, substantially under-performing or, in the worst case, both. To solve the configuration ni ...

Keywords: Disk array, RAID, automatic design

4 [Computing curricula 2001](#)

-  September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press

Full text available:  [pdf\(613.63 KB\)](#)  [html\(2.78 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

5 [Manageability, availability, and performance in porcupine: a highly scalable, cluster-based mail service](#)

-  Yasushi Saito, Brian N. Bershad, Henry M. Levy
August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3


Publisher: ACM Press

Full text available:  [pdf\(2.52 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


This paper describes the motivation, design and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying ...

Keywords: cluster, distributed systems, email, group membership protocol, load balancing, replication

6 [A reliable and scalable striping protocol](#)

-  Hari Adiseshu, Guru Parulkar, George Varghese
August 1996 **ACM SIGCOMM Computer Communication Review , Conference proceedings on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '96**, Volume 26 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(187.15 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Link striping algorithms are often used to overcome transmission bottlenecks in computer networks. Traditional striping algorithms suffer from two major disadvantages. They provide inadequate load sharing in the presence of variable length packets, and may result in non-FIFO delivery of data. We describe a new family of link striping algorithms that solves both problems. Our scheme applies to any layer that can provide multiple

FIFO channels. We deal with variable sized packets ...

7 RAID: high-performance, reliable secondary storage



 Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson
June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(3.60 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...


Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

8 Cluster-based scalable network services



 Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5


Publisher: ACM Press

Full text available:  [pdf\(2.42 MB\)](#)


Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

9 Serverless network file systems



 T. E. Anderson, M. D. Dahlin, J. M. Neefe, D. A. Patterson, D. S. Roselli, R. Y. Wang
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(2.48 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

10 Comparing rebuild algorithms for mirrored and RAID5 disk arrays



 Robert Y. Hou, Yale N. Patt
June 1993 **ACM SIGMOD Record , Proceedings of the 1993 ACM SIGMOD international conference on Management of data SIGMOD '93**, Volume 22 Issue 2

Publisher: ACM Press


Full text available:  [pdf\(945.10 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Several disk array architectures have been proposed to provide high throughput for transaction processing applications. When a single disk in a redundant array fails, the array continues to operate, albeit in a degraded mode with a corresponding reduction in performance. In addition, the lost data must be rebuilt to a spare disk in a timely manner to reduce the probability of permanent data loss. Several researchers have proposed and examined algorithms for rebuilding the failed dis ...


11 Serverless network file systems



 Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S. Roselli, Randolph Y. Wang

February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(2.69 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We propose a new paradigm for network file system design: serverless network file systems. While traditional network file systems rely on a central server machine, a serverless system utilizes workstations cooperating as peers to provide all file system services. Any machine in the system can store, cache, or control any block of data. Our approach uses this location independence, in combination with fast local area networks, to provide better performance and scalability th ...


Keywords: RAID, log cleaning, log structured, log-based striping, logging, redundant data storage, scalable performance

12 Phoenix: a low-power fault-tolerant real-time network-attached storage device

 Anindya Neogi, Ashish Raniwala, Tzi-cker Chiueh


October 1999 **Proceedings of the seventh ACM international conference on Multimedia (Part 1)**

Publisher: ACM Press

Full text available:  [pdf\(1.38 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


Phoenix is a real-time network-attached storage device (NASD) that guarantees real-time data delivery to network clients even across single disk failure. The service interfaces that Phoenix provides are best-effort/real-time reads/writes based on unique object identifiers and block offsets. Data retrieval from Phoenix can be serviced in server push or client pull modes. Phoenix's real-time disk subsystem performance results from a standard cycle-based scan-order disk scheduling mechanism. H ...

13 Query evaluation techniques for large databases

 Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2


Publisher: ACM Press

Full text available:  [pdf\(9.37 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...



Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

14 Declassified disk array architectures with optimal and near-optimal parallelism

 Guillermo A. Alvarez, Walter A. Burkhard, Larry J. Stockmeyer, Flaviu Cristian

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3

Publisher: IEEE Computer Society, ACM Press

Full text available:  [pdf\(1.35 MB\)](#)  Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)
[Publisher Site](#)

This paper investigates the placement of data and parity on redundant disk arrays. Declustered organizations have been traditionally used to achieve fast reconstruction of a failed disk's contents. In previous work, Holland and Gibson identified six desirable properties for ideal layouts; however, no declustered layout satisfying all properties has been published in the literature. We present a complete, constructive characterization of the collection of ideal declustered layouts possessing all ...

15 Q focus: enterprise distributed computing: Enterprise grid computing



 Paul Strong
 July 2005 **Queue**, Volume 3 Issue 6

Publisher: ACM Press

Full text available:  [pdf\(413.13 KB\)](#)  [htm\(31.68 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Grid computing holds great promise for the enterprise data center, but many technical and operational hurdles remain.

16 A transputer T9000 family based architecture for parallel database machines



 Qiang Li, Naphtali Rishe
 December 1993 **ACM SIGARCH Computer Architecture News**, Volume 21 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(638.92 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Parallel computing is a promising way to achieve high performance in a database system. The disk access speed has been a well known bottleneck for database machines, and the data intensive nature and the random communication patterns of databases make the interconnection network in a database machine difficult to design. This article describes the design of a highly parallel, high throughput database machine based on the new T9000 transputer family and a large number of relatively inexpensive dis ...

Keywords: disk array, interconnection network, parallel database machine, semantic model, transputers

17 Data partitioning and load balancing in parallel disk systems



Peter Scheuermann, Gerhard Weikum, Peter Zabback
 February 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(310.27 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Parallel disk systems provide opportunities for exploiting I/O parallelism in two possible ways, namely via inter-request and intra-request parallelism. In this paper, we discuss the main issues in performance tuning of such systems, namely striping and load balancing, and show their relationship to response time and throughput. We outline the main components of an intelligent, self-reliant file system that aims to optimize striping by taking into account the requirements of the applications, an ...

Keywords: Data allocation, Disk cooling, File striping, Load balancing, Parallel disk systems, Performance tuning

18

Fast concurrent access to parallel disks



Peter Sanders, Sebastian Egner, Jan Korst

February 2000 **Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms**

Publisher: Society for Industrial and Applied Mathematics

Full text available:  pdf(986.39 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

19 [Fast detection of communication patterns in distributed executions](#)

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Publisher: IBM Press

Full text available:  pdf(4.21 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

20 [Asynchronous scheduling of redundant disk arrays](#)

Peter Sanders

July 2000 **Proceedings of the twelfth annual ACM symposium on Parallel algorithms and architectures**

Publisher: ACM Press

Full text available:  pdf(161.35 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)





Random redundant allocation of data to parallel disk arrays can be exploited to achieve low access delays. New algorithms are proposed which improve the previously known shortest queue algorithm by systematically exploiting that scheduling decisions can be deferred until a block access is actually started on a disk. These algorithms are also generalized for coding schemes with low redundancy. Using extensive experiments, practically important quantities are measured which have so far eluded ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

Search Results[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)Results for "((raid<and>(logical<near/1>disk)<and> (logical<near/1>entities)<and>(l..." [e-mail](#)Your search matched **0** documents.A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

[View Session History](#)[Modify Search](#)[New Search](#) [»](#)☐ Check to search only within this results set

» Key

Display Format: ☒ Citation ☐ Citation & Abstract

IEEE JNL IEEE Journal or Magazine

IEEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistance with your search.

Indexed by
 Inspec®[Help](#) [Contact Us](#) [Privacy & Policy](#)

© Copyright 2005 IEEE - All rights reserved.

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	0	((logical adj disk) with (logical adj entities) with copying).clm.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 13:51
L2	0	((logical adj disk) with (logical adj entities) same copying).clm.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/09 13:51